

Spatial Decompositions for Large Scale SVMs

Philipp Thomann

Institut für Stochastik und Anwendungen

Universität Stuttgart

Pfaffenwaldring 57, 70569 Stuttgart

`philipp.thomann@mathematik.uni-stuttgart.de`

Ingo Steinwart

Institut für Stochastik und Anwendungen

Universität Stuttgart

Pfaffenwaldring 57, 70569 Stuttgart

`ingo.steinwart@mathematik.uni-stuttgart.de`

Ingrid Blaschzyk

Institut für Stochastik und Anwendungen

Universität Stuttgart

Pfaffenwaldring 57, 70569 Stuttgart

`ingrid.blaschzyk@mathematik.uni-stuttgart.de`

Mona Meister

Robert Bosch GmbH

70442 Stuttgart

`mona.meister@de.bosch.com`

December 4, 2016

Abstract

Although support vector machines (SVMs) are theoretically well understood, their underlying optimization problem becomes very expensive if, for example, hundreds of thousands of samples and a non-linear kernel are considered. Several approaches have been proposed in the past to address this serious limitation. In this work we investigate a decomposition strategy that learns on small, spatially defined data chunks. Our contributions are two fold: On the theoretical side we establish an oracle inequality for the overall learning method using the hinge loss, and show that the resulting rates match those known for SVMs solving the complete optimization problem with Gaussian kernels. On the practical side we compare our approach to learning SVMs on small, randomly chosen chunks. Here it turns out that for comparable training times our approach is significantly faster during testing and also reduces the test error in most cases significantly. Furthermore, we show that our approach easily scales up to 10 million training samples: including hyper-parameter selection using cross validation, the entire training only takes a few hours on a single machine. Finally, we report an experiment on 32 million training samples.

1 Introduction

Kernel methods are thoroughly understood from a theoretical perspective, used in many settings, and are known to give good performance for small and medium sized training sets. Yet they suffer from their computational complexity that grows quadratically in space and at least quadratically in time. For example, storing the entire kernel matrix to avoid costly recomputations in e.g. 64GB of memory, one can consider at most 100 000 data points. In the last 15 years there has been wide research to circumvent this barrier (cf. [2] for an overview): Sequential Minimal Optimization [12] allows for caching

kernel rows so that the memory barrier is lifted, parallel solvers try to leverage from recent advances in hardware design, matrix approximations [20] replace the original $n \times n$ kernel matrix, where n is the number of samples, by a smaller approximation, the random Fourier feature method [13] approximates the kernel function by another kernel having an explicit and low-dimensional feature map, which, in second step, makes it possible to solve the primal problem instead of the usually considered dual problem, see [15] for a recent theoretical investigation establishing optimal rates for this approximation. Moreover, iterative strategies [14, 11] modify the underlying regularization approach, sometimes even by controlled early stopping. Finally, various decomposition strategies have been proposed, which, in a nutshell, solve many small rather than one large optimization problem.

In this work we also focus on a data decomposition strategy. Such strategies have been investigated at least since [3, 19], and probably the most simple such strategy, called *random chunks*, splits the data by random into chunks of some given size, train for each chunk, and finally average the decision functions obtained on every chunk to one final decision function. From a theoretical perspective this approach has been investigated recently in [23].

Obviously, the drawback with this approach is that the test sample has to be evaluated on every chunk: for example, if every chunk uses 80% of its training samples as support vectors, then the test sample has to be evaluated using 80% of the entire training set. From this perspective, a more interesting strategy is hence to decompose the data set into spatially defined cells, because in this case every test sample has only to be evaluated using the cell it belongs to. In our example above, this amounts to 80% of the cell size, instead of 80% of the entire training set size. Clearly, the difference of both costs can be very significant if, for example the cell size ranges in between say a few thousands, but the training set contains millions of examples. The natural next question is of course, whether and by how much one suffers from this approach in terms of test accuracy.

Spatial decompositions for SVMs can be obtained in many different ways, e.g., by clusters [6, 7], decision trees [1, 21, 5], and k -nearest neighbors [22]. Most of these strategies are investigated experimentally, yet only very few theoretical results are known: [10] proves universal consistency of localized versions of SVMs. Oracle inequalities and optimal learning rates for least squares SVMs have been shown in [9].

In the theoretical part of this work we expand the results in [9] for the least squares case to the hinge loss. Besides the obviously different treatment of the approximation error, the main difference in our work is that the least squares loss allows to use an optimal variance bound, whereas for the hinge loss such an optimal variance bound is only available for distributions satisfying the best version of Tsybakov’s noise condition, see (2). In general, however, only a weaker variance bound is possible, which in turn makes the technical treatment harder and, surprisingly, the conditions on the cell radii that guarantee the best known rates, more restrictive.

In the experimental part we have trained local SVMs, that is, SVMs that are trained on disjunct spatially defined cells, on some well-known data sets to demonstrate that these local SVMs are an efficient way to tackle large-scale problems with millions of samples. Because we use full 5-fold cross validation on a 10×10 hyper-parameter grid this shows that SVMs promise to achieve fully automatic machine learning. Still it takes only 2-6 hours to train on data sets with 10 millions of samples and 28 features, using only 10-30GB of memory and a single computer. Finally, we also distributed one training experiment onto 11 machines to show how it performs for a data set with 30 million samples and 631 features.

In Section 2 we give a description of the local SVM approach. In Section 3 we define local SVMs for the hinge loss and Gaussian kernels and state the theoretical results. The proofs can be found in the supplement. To motivate the theoretical results we compare these results in Section 4 with the one obtained by a toy example. In Section 5 we describe the extensive experiments we performed.

2 The local SVM approach

In this section we briefly describe the local SVM approach. To this end, let $D := ((x_1, y_1), \dots, (x_n, y_n))$ be a data set of length n , where $x_i \in \mathbb{R}^d$ and $y_i \in \mathbb{R}$. Local SVMs construct a function f_D by solving SVMs on many spatially defined small chunks. To be more precise, let $(A_j)_{j=1, \dots, m}$ be an arbitrary

partition of the input space X . We define for every $j \in \{1, \dots, m\}$ the local data set D_j by

$$D_j := \{(x, y) \in D : x \in A_j\}.$$

Then, one learns an individual SVM on *each* cell by solving for a regularization parameter $\lambda > 0$ the optimization problem

$$f_{D_j, \lambda} = \arg \min_{f \in H_j} \lambda \|f\|_{H_j}^2 + \frac{1}{n} \sum_{x_i, y_i \in D_j} L(y_i, f(x_i))$$

for every $j \in \{1, \dots, m\}$, where H_j is a reproducing kernel Hilbert space over A_j with reproducing kernel $k_j : A_j \times A_j \rightarrow \mathbb{R}$, see [16, Chapter 4], and where $L : Y \times \mathbb{R} \rightarrow [0, \infty)$ is a function describing our learning goal. The final decision function $f_{D, \lambda} : X \rightarrow \mathbb{R}$ is then defined by

$$f_{D, \lambda}(x) := \sum_{j=1}^m \mathbf{1}_{A_j}(x) f_{D_j, \lambda}(x).$$

To illustrate the advantages of this approach, let us assume that the size of the data sets D_j is nearly the same for all $j \in \{1, \dots, m\}$, that means $|D_j| \approx n/m$. For example, if the data is uniformly distributed, it is not hard to show that the latter assumption holds, and if the data lies uniformly on a low-dimensional manifold, the same is true, since empty cells can be completely ignored. Now, the calculation of the kernel matrices $K_j^{i, l} = k_j(x_i, x_l)$ for $x_i, x_l \in D_j$ scales as

$$O\left(m \cdot \left(\frac{n}{m}\right)^2 \cdot d\right) = O\left(\frac{n^2 \cdot d}{m}\right).$$

In comparison, for a global SVM the calculation of the kernel matrix $K^{i, l} = k(x_i, x_l)$ scales as

$$O(n^2 \cdot d),$$

such that splitting and multi-thread improve that scaling by $1/m$. Similarly, it is well-known that the time complexity of the solver is

$$O\left(m \cdot \left(\frac{n}{m}\right)^{1+a}\right) = O\left(n \cdot \left(\frac{n}{m}\right)^a\right), \quad (1)$$

where $a \in [1, 2]$ is a constant, and the test time scales as the kernel calculation (see Table 6 in the supplement for experimental corroboration). In all three phases we thus see a clear improvement over a globally trained SVM. Moreover, while SVMs trained on random chunks have the same complexities for the kernel matrix and the solver, they only have the bad complexity of the global approach during testing.

3 Oracle inequality and learning rates for local SVMs with hinge loss

The aim of this section is to theoretically investigate the local SVM approach for binary classification. To this end, we define for a measurable function $L : Y \times \mathbb{R} \rightarrow [0, \infty)$, called loss function, the L -risk of a measurable function $f : X \rightarrow \mathbb{R}$ by

$$\mathcal{R}_{L, P}(f) = \int_{X \times Y} L(y, f(x)) dP(x, y).$$

Moreover, we define the optimal L -risk, called the Bayes risk with respect to P and L , by

$$\mathcal{R}_{L, P}^* := \inf \{ \mathcal{R}_{L, P}(f) \mid f : X \rightarrow \mathbb{R} \text{ measurable} \}$$

and call a function $f_{L, P}^* : X \rightarrow \mathbb{R}$, attaining the infimum, Bayes decision function. Given a data set $D := ((x_1, y_1), \dots, (x_n, y_n))$ drawn in an i.i.d. fashion from a probability measure P on $X \times Y$,

where $X \subset \mathbb{R}^d$ and $Y := \{-1, 1\}$, the learning target in binary classification is to find a decision function $f_D : X \rightarrow \mathbb{R}$ such that $\text{sign } f_D(x) = y$ for new data (x, y) with high probability. A loss function describing our learning goal is the classification loss $L_{\text{class}} : Y \times \mathbb{R} \rightarrow [0, \infty)$, defined by

$$L_{\text{class}}(y, t) := \mathbf{1}_{(0, \infty]}(y \text{ sign } t),$$

where $\text{sign } 0 := 1$. Another possible loss function is the hinge loss $L_{\text{hinge}} : Y \times \mathbb{R} \rightarrow [0, \infty)$, defined by

$$L_{\text{hinge}}(y, t) := \max\{0, 1 - yt\}$$

for $y = \pm 1, t \in \mathbb{R}$. Since the hinge loss is convex and therefore suitable for SVMs and since a well-known result by Zhang, e.g., [16, Theorem 2.31] shows that

$$\mathcal{R}_{L_{\text{class}}, P}(f) - \mathcal{R}_{L_{\text{class}}, P}^* \leq \mathcal{R}_{L_{\text{hinge}}, P}(f) - \mathcal{R}_{L_{\text{hinge}}, P}^*$$

for all functions $f : X \rightarrow \mathbb{R}$, we consider in the following the hinge loss in our theory and write $L := L_{\text{hinge}}$. We remark that for the hinge loss it suffices to consider the risk for functions values restricted to the interval $[-1, 1]$, since this does not worsen the loss and thus the risk. Therefore, we define by

$$\hat{t} := \max\{-1, \min\{t, 1\}\}$$

for $t \in \mathbb{R}$ the clipping operator, which restricts values of t to $[-1, 1]$, see [16, Chapter 2.2]. In order to derive a bound on the excess risks $\mathcal{R}_{L, P}(\hat{f}_{D, \lambda}) - \mathcal{R}_{L, P}^*$, and in order to derive learning rates, we recall some notions from [16, Chapter 8], which describe the behaviour of P in the vicinity of the decision boundary. To this end, let $\eta : X \rightarrow [0, 1]$ be a version of the posterior probability of P , that means that the probability measures $P(\cdot | x)$, defined by $P(y = 1 | x) =: \eta(x)$, $x \in X$, form a regular conditional probability of P . We write

$$\begin{aligned} X_1 &:= \{x \in X : \eta(x) > 1/2\}, \\ X_{-1} &:= \{x \in X : \eta(x) < 1/2\}. \end{aligned}$$

Then, the function $\Delta_\eta : X \rightarrow [0, \infty)$ defined by

$$\Delta_\eta(x) := \begin{cases} \text{dist}(x, X_1) & \text{if } x \in X_{-1}, \\ \text{dist}(x, X_{-1}) & \text{if } x \in X_1, \\ 0 & \text{otherwise,} \end{cases}$$

is called distance to the decision boundary, where $\text{dist}(x, A) := \inf_{x' \in A} \|x - x'\|_2$. On the basis of this definition we can describe the mass of the marginal distribution P_X of P around the decision boundary by the following exponents: We say that P has margin-noise exponent (MNE) $\beta \in (0, \infty)$ if there exists a constant $c_{\text{MNE}} \geq 1$ such that

$$\int_{\{\Delta_\eta(x) < t\}} |2\eta(x) - 1| dP_X(x) \leq c_{\text{MNE}} t^\beta$$

for all $t > 0$. That is, the margin-noise exponent β measures the mass and the noise, i.e. points $x \in X$ with $\eta(x) \approx 1/2$, around the decision boundary. Hence, we have a large margin-noise exponent if we have low mass and/or high noise around the decision boundary. Furthermore, we say that P has noise exponent (NE) $q \in [0, \infty]$ if there exists a constant $c_{\text{NE}} > 0$ such that

$$P_X(\{x \in X : |2\eta(x) - 1| < \varepsilon\}) \leq (c_{\text{NE}})^q \quad (2)$$

for all $\varepsilon > 0$. Thus, the noise exponent q , which corresponds to Tsybakov's noise condition, introduced in [18], measures the amount of noise, but does not locate it. For examples of typical values of these exponents and relations between them we refer the reader to [16, Chapter 8].

For the theory of local SVMs it is necessary to specify the partition. To this end, we assume in the following that $X \subset [-1, 1]^d$ and we define the radius r_T of a set $T \subset X$ by

$$r_T = \inf\{\varepsilon > 0 : \exists s \in T \text{ such that } T \subset B_2(s, \varepsilon)\},$$

where $B_2(s, \varepsilon) := \{t \in T : \|t - s\|_2 \leq \varepsilon\}$ with Euclidean norm $\|\cdot\|_2$ in \mathbb{R}^d . In the following let $(A_j)_{j=1, \dots, m}$ be a partition of X and $r > 0$ such that for every $j \in \{1, \dots, m\}$ we have $\mathring{A}_j \neq \emptyset$ and

$$r_{A_j} < r \leq 16m^{-\frac{1}{d}}. \quad (3)$$

A simple partition that fulfills the latter condition is for example the partition by cubes with a specific side length. We refer the reader to Section 4 for the computation of another type of partition satisfying (3).

In the following we restrict ourselves to local SVMs with Gaussian kernels. For this purpose, we denote for every $j \in \{1, \dots, m\}$ by $H_{\gamma_j}(A_j)$ the RKHS over A_j with Gaussian kernel $k_{\gamma_j} : A_j \times A_j \rightarrow \mathbb{R}$, defined by

$$k_{\gamma_j}(x, x') := \exp(-\gamma_j^{-2} \|x - x'\|_2^2)$$

for some width $\gamma_j > 0$. Furthermore, we define for $f \in H_{\gamma_j}(A_j)$ the function $\hat{f} : X \rightarrow \mathbb{R}$ by

$$\hat{f}(x) := \begin{cases} f(x), & x \in A_j, \\ 0, & x \notin A_j. \end{cases}$$

Then, according to [9, Lemma 2] the space $\hat{H}_{\gamma_j} := \{\hat{f} : f \in H_{\gamma_j}(A_j)\}$ equipped with the norm

$$\|\hat{f}\|_{\hat{H}_{\gamma_j}} := \|f\|_{H_{\gamma_j}(A_j)}$$

is an RKHS over X . Hence, local SVMs for Gaussian kernels solve for every $j \in \{1, \dots, m\}$ the optimization problem

$$f_{D_j, \lambda_j, \gamma_j} = \arg \min_{\hat{f} \in \hat{H}_{\gamma_j}} \lambda_j \|\hat{f}\|_{\hat{H}_{\gamma_j}}^2 + \frac{1}{n} \sum_{x_i, y_i \in D_j} L(y_i, f(x_i)),$$

where $\lambda_j, \gamma_j > 0$. Then, for the vectors $\gamma := (\gamma_1, \dots, \gamma_m) > 0$ and $\lambda := (\lambda_1, \dots, \lambda_m) > 0$ the decision function $f_{D, \lambda, \gamma} : X \rightarrow \mathbb{R}$ is defined by

$$f_{D, \lambda, \gamma}(x) := \sum_{j=1}^m f_{D_j, \lambda_j, \gamma_j}(x). \quad (4)$$

Note that the clipped decision function $\hat{f}_{D, \lambda, \gamma} : X \rightarrow [-1, 1]$ is then defined by the sum of the clipped empirical solutions $\hat{f}_{D_j, \lambda_j, \gamma_j}$ since for all $x \in X$ there is exactly one $f_{D_j, \lambda_j, \gamma_j}$ with $f_{D_j, \lambda_j, \gamma_j}(x) \neq 0$. We finally introduce the following set of assumptions:

- (A) Let $(A_j)_{j=1, \dots, m}$ be a partition of X and $r > 0$, such that for every $j \in \{1, \dots, m\}$ we have $\mathring{A}_j \neq \emptyset$ and (3). In addition, for every $j \in \{1, \dots, m\}$ let $H_{\gamma_j}(A_j)$ be the RKHS of the Gaussian kernel k_{γ_j} over A_j with width $\gamma_j > 0$. Furthermore, we use the notation $\gamma_{\max} := \max\{\gamma_1, \dots, \gamma_m\}$.

Based on the assumptions made above, we now present an upper bound on the excess risk for the hinge loss and Gaussian kernels.

Theorem 3.1. *Let $Y = \{-1, 1\}$ and let $L : Y \times \mathbb{R} \rightarrow [0, \infty)$ be the hinge loss. Let P be a distribution on $X \times Y$ that has margin-noise exponent $\beta \in (0, \infty)$ and noise exponent $q \in [0, \infty]$. Furthermore let (A) be satisfied. Then, for all $p \in (0, 1)$, $n \geq 1$, $\tau \geq 1$ with $\tau \leq n$, $\lambda = (\lambda_1, \dots, \lambda_m) \in (0, \infty)^m$, $\gamma = (\gamma_1, \dots, \gamma_m) \in (0, r]^m$ the SVM given by (4) satisfies*

$$\begin{aligned} \mathcal{R}_{L, P}(\hat{f}_{D, \lambda, \gamma}) - \mathcal{R}_{L, P}^* &\leq C_{\beta, d, p, q} \left(\sum_{j=1}^m \lambda_j \gamma_j^{-d} + \gamma_{\max}^\beta + \left(\frac{\tau}{n}\right)^{\frac{q+1}{q+2}} \right. \\ &\quad \left. + \left(r^{2p} \left(\sum_{j=1}^m \lambda_j^{-1} \gamma_j^{-\frac{d+2p}{p}} P_X(A_j) \right)^p n^{-1} \right)^{\frac{q+1}{q+2-p}} \right) \end{aligned}$$

with probability P^n not less than $1 - 3e^{-\tau}$, where $C_{\beta, d, p, q} > 0$ is a constant depending only on d, β, p and q .

The main idea used in the proof is that $f_{D,\lambda,\gamma}$ is an SVM solution for a particular RKHS. To be more precise, it is easy to show with a generalization of [9, Lemma 3] that for RKHSs \hat{H}_{γ_j} on X and a vector $\lambda := (\lambda_1, \dots, \lambda_m) > 0$ the direct sum

$$H := \bigoplus_{j=1}^m \hat{H}_{\gamma_j} = \left\{ f = \sum_{j=1}^m f_j : f_j \in \hat{H}_{\gamma_j} \text{ for all } j \in J \right\}$$

is again an RKHS equipped with the norm

$$\|f\|_H^2 = \sum_{j=1}^m \lambda_j \|f_j\|_{\hat{H}_{\gamma_j}}^2.$$

Obviously, $f_{D,\lambda,\gamma} \in H$. We remark that the latter construction actually holds for RKHSs with arbitrary kernels. Besides that it is necessary for the proof to choose suitable RKHSs in order to bound the approximation error, it is also necessary to bound the entropy numbers of the operator $\text{id} : H_{\gamma_j}(A_j) \mapsto L_2(P_{X|A_j})$, which describe in some sense the size of the RKHS $H_{\gamma_j}(A_j)$. Unfortunately, such bounds contain constants which depend on the cells A_j and which are in general hard to control. However, for Gaussian kernels we obtain such bounds if we restrict γ_j by r for every $j \in \{1, \dots, m\}$, which explains our restriction to Gaussian kernels. Moreover, the proof actually shows that $\gamma_j \leq r$ can be replaced by $\gamma_j \leq cr$ for a constant c , which is independent of p, m, τ, λ and γ . Furthermore, using this entropy number bound leads to a dependence of a parameter p on the right-hand side of the oracle inequality, where small p lead to an unknown behaviour in the constant $C_{\beta,d,p,q}$. This problem is well-known in the Gaussian kernel case, see [17] or [8], and there is not given a solution for this problem yet.

Let us assume that all cells have the same kernel parameter γ and regularization parameter λ . Then, the oracle inequality stated in Theorem 3.1 coincides up to constants and the parameter p , which can be chosen arbitrary small, the oracle inequality for global SVMs stated in [16, Theorem 8.25] for $\tau = \infty$. Furthermore, we remark that our oracle inequality is formally similar to the inequality for local SVMs for the least square loss, see [9, Theorem 7] if we assume that the Bayes decision function is contained in a Besov space with smoothness $\alpha = \beta/2$.

In the next theorem we show learning rates by choosing appropriate sequences of $r_n, \lambda_{n,j}$ and $\gamma_{n,j}$.

Theorem 3.2. *Let $\tau \geq 1$ be fixed and $\nu \in \left(0, \frac{q+1}{\beta(q+2)+d(q+1)}\right]$. Under the assumptions of Theorem 3.1 and with*

$$\begin{aligned} r_n &= c_1 n^{-\nu}, \\ \lambda_{n,j} &= c_2 r_n^d n^{-\frac{(\beta+d)(q+1)}{\beta(q+2)+d(q+1)}}, \\ \gamma_{n,j} &= c_3 n^{-\frac{(q+1)}{\beta(q+2)+d(q+1)}}, \end{aligned}$$

for every $j \in \{1, \dots, m_n\}$, we have for all $n \geq 1$ and $\xi > 0$ that

$$\mathcal{R}_{L,P}(\hat{f}_{D,\lambda_n,\gamma_n}) - \mathcal{R}_{L,P}^* \leq C_{\beta,\nu,\xi,d,q} \tau^{\frac{q+1}{q+2}} \cdot n^{-\frac{\beta(q+1)}{\beta(q+2)+d(q+1)} + \xi}$$

holds with probability P^n not less than $1 - 3e^{-\tau}$, where $\lambda_n := (\lambda_{n,i})_{i=1,\dots,m_n}$ as well as $\gamma_n := (\gamma_{n,i})_{i=1,\dots,m_n}$ and where $C_{\beta,\nu,\xi,d,q}, c_1, c_2, c_3$ are positive constants.

Note that the restriction of the parameter ν in Theorem 3.2 is set to ensure that $\sup_{n \geq 1} \gamma_{n,j}/r_n < \infty$, as we mentioned after Theorem 3.1. The learning rate stated in Theorem 3.2 coincides always with the fastest known rate which can be achieved by a global SVM, cf. [16, Theorem 8.26 and (8.18)]. Let us consider some special cases for the parameters β and q . In the case of "benign noise", that is $q = \infty$, our learning rate reduces to

$$n^{-\frac{\beta}{\beta+d} + \xi},$$

and is only less sensitive to the dimension d if β is large. Next, let us assume that $q = 1$ such that P has a rather moderate noise concentration and that $\beta = 2$, which means that we have additionally

much mass around the decision boundary. For examples of distributions having these parameters we refer to the examples in [16, Chapter 8]. The chosen parameters q and β yield the rate

$$n^{-\frac{2}{3+d}+\xi}$$

and we observe that the dimension d has a high impact, which means that the rate gets worse the higher d . We refer the reader to Section 4, where we created a toy example for a distribution having these parameters β and q and where we compare this rate with the experimental one. Since the class of considered distributions contains the ones, whose marginal distribution P_X is the uniform distribution, a bad dependence on d is not surprisingly as these distributions usually among those which cause the curse of dimensionality. However, if the data lies for example on a $d_{\mathcal{M}}$ -(low)-dimensional rectifiable manifold \mathcal{M} , we believe that one is able to improve the rate since we then have to learn the local SVM only on a few cells. That is, instead of learning on $m = O(r^d)$ cells, we only have to learn on $m = O(r^{d_{\mathcal{M}}})$ cells.

Clearly, to obtain the rates in Theorem 3.2 we need to know the parameters β and q . However, such rates can also be obtained by a data-dependent parameter selection strategy without knowing the parameters. For example, [9] presented for the least-square loss the training validation Voronoi partition support vector machine (TV-VP-SVM) and showed that this learning method achieves adaptively the same rates as the rates for local SVM for the least-square loss. We remark that this method can be adapted to our case, that is for the hinge loss, since a key ingredient is an oracle inequality having the structure of [9, Theorem 7], which is given in our case, as we mentioned in the discussion before Theorem 3.2. For more details to parameter selection methods we refer the reader to [9, Section 5] and [16, Chapter 6.5 and Chapter 8.2]. At this point we finally remark that the mentioned adaptive strategies obtain the rate

$$n^{-\frac{\beta(q+1)}{\beta(q+2)+d(q+1)}+\xi}$$

for all β and q satisfying $\nu \leq \frac{\beta(q+1)}{\beta(q+2)+d(q+1)}$. Consequently, for larger ν , which lead to faster training times, the range of adaptivity becomes smaller, and vice-versa.

4 Toy Example

To illustrate the theoretical results we now give a toy example by mixing two multivariate Gaussians, one for $y = 1$ and one for $y = -1$, see Figure 1. We define $X := [-2, 2]^d$ and $x_1 := (1, 0, \dots, 0) \in X$




Figure 1: Toy model: Mixture of two Gaussians. The decision boundary is depicted as almost a circle and the dotted lines mark a Voronoi partition induced by the black center samples.

and $\vartheta := 0.6$. Moreover, we set for all events $A \subset X \times Y$:

$$P(A) := \vartheta \int_{A^+} \phi(x) dx + (1 - \vartheta) \int_{A^-} \psi(x) dx$$

where $A^+ := \{x \mid (x, 1) \in A\}$ and $A^- := \{x \mid (x, -1) \in A\}$ and where ϕ and ψ are the densities of the multivariate normal distribution around the origin, and x_1 resp. and of variance 1 and $\frac{1}{8}$ resp. on X . In this case $\eta(x)$ can be calculated and it is easy to see that the decision boundary is an ellipsoid, the noise exponent is $q = 1$, and the margin-noise exponent is $\beta = 2$. Hence, we can use Theorem 3.2 by using

$$\begin{aligned} \nu &= \frac{1}{3+d}, & r_n &= c_1 n^{-\frac{1}{3+d}}, \\ \lambda_n &= c_2 n^{-\frac{2+d}{3+d}}, & \gamma_n &= c_3 n^{-\frac{1}{3+d}}. \end{aligned}$$

to achieve the rate $n^{-\frac{2}{3+d}+\varepsilon}$. a (even faster) rate than the theoretical from Theorem 3.2.

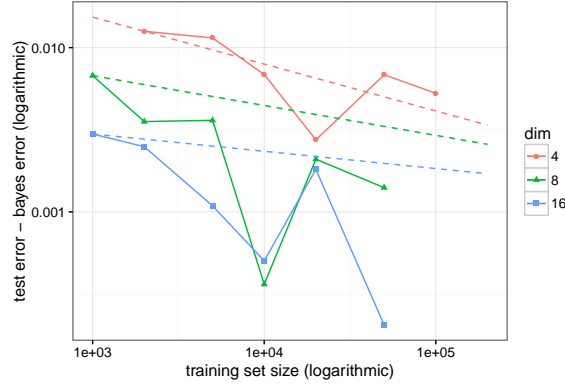


Figure 2: Obtained excess risks for the toy models at dimensions 4, 8, 16 where we decrease the maximal data radius per cell as in Theorem 3.2. The theoretical rates from Theorem 3.2 are marked in dashed lines with normalization to match the first data point. We only performed one experiment per model, but averaging over more runs should smoothen out the kinks.

5 Experiments

We now use the developed methods for large scale data sets in order to understand whether the theoretical and synthetic results also transfer to real world data. We intend to demonstrate that partitioning allows for large n to give efficiently good results, but only if one uses spatial decomposition. As discussed above, global kernel methods become unfeasible for $n \geq 100\,000$, but there is already known the decomposition method of *random chunks*: split the data into samples and train on these smaller sets first, then average predictions over these samples. We will see that spatial decompositions often outperform random chunks in terms of test error—while testing is much faster. We did not consider any other method for speeding up SVMs, e.g. random fourier features, since these can be naturally combined with our spatial approach. Indeed, if such a method is faster than a global SVM on data sets of, say 20.000 and more samples, then one could also use this method on each cell of that size, which in turn would further decrease the overall training time of our approach. In this sense, our reported training times are a kind of worst-case scenario.

In Section 5.1 we will demonstrate that the spatial decomposition approach even can be used for problems with dozens of millions of samples and hundreds of dimensions, that is too big for a single machine.

For the data sets, we looked on the UCI and LIBSVM repositories for the biggest examples, which did have firstly a suitable learning target, secondly only numerical features, and thirdly not too many dimensions since for high dimensions spatial decompositions become at least debatable. Hence we did




Figure 3: Test errors for the large-scale data sets. The training set size is the most influential factor for the spatial decompositions. Bigger cell sizes give only sometimes earlier better results. Random chunk decomposition only profits from bigger chunks and not from more chunks.

not consider any image data sets. See Table 1 for an overview of the data sets and the train/test splits. We trained on one hand using the full training set, and on the other hand using a sample of the full training set of the 6 sizes $n = 50\,000, \dots, 1\,000\,000, 2\,500\,000$.

Even though our theoretical results are by control of radii, in applications it is more important to control the computational cost, namely the maximal number of samples in the cells as we discussed in Section 2. Hence, for each cell size 2000, 5000, 10000, 15000, each of the 7 training sets was first split spatially into cells of that size and a SVM was trained on each cell. This means that on each cell five-fold cross validation was performed for a 10×10 geometrically spaced hyper-parameter grid where the endpoints were scaled to accommodate the number of samples in every fold, the cell size, and the dimension. Then, for the full testing set the test error was computed by assigning each test sample to the cell it spatially belongs to and the cell’s SVM was used to predict a label for it. For comparison we also performed these experiments using random chunks: Each training set was split uniformly into chunks of size 2000, 5000, 10000, 15000. There, due to time constraints, for the bigger data sets we calculated the testing error using a test set of size 100 000.

We used our own SMO-type implementation in C++, which we will make available. As architecture, we used Intel® Xeon® CPUs (E5-2640 0 at 2.50GHz, May 2013) with Ubuntu linux. There were two NUMA-sockets each with a CPU having 6 physical cores, but multi-threading was only used to compute the kernel matrix in training and the test error. The partitioning of the data and the solver are single-threaded. This hence could be enhanced significantly in further development. Even though there were 128GB RAM per NUMA-socket available, the processes were limited to use at most 64GB

Table 1: Overview of the data sets we used.

Name	train size	test size	dims	Source
HIGGS	9 899 999	1 100 001	28	UCI
HEPMASS	7 000 000	3 500 000	28	UCI
GASSENSOR	7 548 087	838 678	18	UCI
SUSY	4 499 999	50 0001	18	UCI
COVTYPE	464 429	116 583	54	UCI
COD-RNA	231 805	99 347	9	LIBSVM
SKIN	196 045	49 012	3	UCI

to give results which could be compared on other workstations. The smaller cell sizes use considerably less memory. But this way, we restricted us to only use cell sizes up to 15 000. For random chunks in the biggest data sets even that was not always possible.

Table 2: Time to train (in minutes) the local SVM including 5-fold cross-validation on a 10×10 hyper-parameter grid.

	2000	5000	10000	15000
	training time			
HIGGS	308	679	1358	1992
HEPMASS	145	316	624	964
GASSENSOR	90	214	421	636
SUSY	121	261	513	779
COVTYPE	9	18	39	54
COD-RNA	4	9	16	25
SKIN	2	6	10	16

Table 3: Time to predict using random chunks divided by time to predict using spatial decomposition. Since RC becomes to expensive this is using 100 000 training and test samples.

	2000	5000	10000	15000
HIGGS	74	34	15	10
HEPMASS	84	28	14	6
GASSENSOR	667	513	254	414
SUSY	88	35	19	11
COVTYPE	153	61	26	19
COD-RNA	139	52	18	6
SKIN	25	10	8	4

The results on real world data sets give a clear picture. It can be seen in Figure 3 that increasing the training set size enhances the test error in most cases dramatically. All data sets but SUSY show that by using spatial partitioning one is able to attack large-scale problems nicely. Remark that all cell sizes give almost the same error. Yet smaller cell sizes give it faster (see Table 2), only for HIGGS and HEPMASS bigger cells achieve the same test error already using fewer training data. Bigger cell sizes can give better test error, although in the trade-off time vs. error, they play not too big of a role, cf. Figure 4 in the supplement. In contrast, for the random chunks strategy the cell size plays the most crucial role for the test error and the error quickly saturates at a some value. The error can only be made smaller by using bigger cell sizes, and hence much more expensive training and even more testing time, see Table 3. Finally, our spatial decomposition seems overwhelmed with SUSY and exposes the same saturation effect as in random chunk and does not achieve to beat that.

5.1 Big-data: ECBDL 2014

One of the most important aspects of spatial decompositions is that this makes the process cloud scalable, since the training and testing on the cells trivially can be assigned to different workers, once the data is split. To demonstrate this we used Apache Spark, an in-memory map/reduce framework.

Table 4: More detailed times at fixed cell size 2000 (in seconds). The testing phase is given per 100 000 test samples. The last column gives the maximal resident size over all of training and testing (in GB).

name	part.	time (in seconds) used in phase					RAM
		kernel calc.	solver	valid.	sel	test	
HIGGS	40	2887	11862	994	1030	322	31
HEPMASS	26	2022	4235	682	560	208	22
GASSENSOR	21	2354	729	633	541	166	21
SUSY	18	1296	4419	448	374	141	13
COVTYPE	3	138	206	44	47	13	3
COD-RNA	1	69	76	20	21	7	1
SKIN	1	53	21	18	11	3	1

The data set was saved on a Hadoop distributed file system on one master and up to 11 worker machines of the above type. In a first step, the data was split into coarse cells by the following procedure. A subset of the training data was sampled and sent to the master machine where 1000-2000 centers were found and these centers were sent back to the worker machines. Now each worker machine could assign locally to every of its samples the coarse cell in the Voronoi sense. Finally a Spark-shuffle was performed: Every cell was assigned to one of the workers and all its samples were sent to that worker. This procedure is quite standard for Spark. It had to be performed only once for the data set.

In the second step every such coarse cell—now being on one physical machine—was used for training by our C++ implementation discussed above: this in particular means that each coarse cell was again split into fine cells of some specified size and then the TV-VP-SVM method was used. For these experiments we used a 20×20 hyper-parameter grid. Obviously this now was done in parallel on all worker nodes. The test set was also split into the coarse cells and then by our implementation further into fine cells for prediction.

To give an experimental evaluation we used the classification data set introduced in the *Evolutionary Computation for Big Data and Big Learning Workshop*¹ (ECBDL). The data set considers contact points of polymers. It has 631 dimensions and 32 million training (60GB disk space) and 2.9 million test samples (5.1GB disk space). One of the major challenges with this is a rather strong imbalance in the labels—there are 98% of negative samples. Therefore the competition used a scoring of "TruePositiveRate · TrueNegativeRate".

Certainly this data set is not ideally suited for local SVMs: on one hand it certainly has higher dimensions than we would hope. On the other hand the imbalance in the data has to be treated by hand. For this we used the hinge loss with weight 0.987 after trying out different weights on a validation sample of size 100 000. For the training calculations, the use of the full 128GB memory per socket allowed us to use (fine) cell sizes up to 100 000. The splitting was not optimized and took about an hour. Training and testing took 32.2 hours on our 11 worker machines, see Table 5.

Our off-the-cuff scores range from 0.422 to 0.456. That would have landed us in the middle of the scores at the beginning of the competition. Of the seven teams three had their first submissions between 0.3 and 0.42 and at some point found a way to boost it to 0.45 or more. One of these, the team HYPERENS, used *standard and budgeted SVMs with bayesian optimisation of parameters* and started with submissions scoring around 0.34–0.38 and after ten days found a way to score over 0.45 and achieving in the end 0.489 using *4.7 days of parameter optimisation in a 16-core machine*. The three best teams scored from start over 0.47. The winning team EFDAMIS used feature weighting and random forests and its best model took *39h of Wall-clock time in a 144-core cluster (not used exclusively)* which is comparable to the training time of our best model.

Generally, we suppose that in such competitions the best results are achieved by careful hand-optimisation of the features and the hyper-parameters. We on the other hand aspire to realize fully automatic learning and hence are not aiming to beat such results.

¹ See <http://cruncher.ncl.ac.uk/bdcomp/> and <http://cruncher.ncl.ac.uk/bdcomp/BDCOMP-final.pdf>.

Table 5: Results for the ECBDL’14 data set. Increasing the cell size made it possible to increase the score.

cell size	time	Score	TPR	TNR	work nodes
10 000	4.8	0.422	0.656	0.643	7
15 000	7.2	0.433	0.664	0.652	7
20 000	9.3	0.438	0.664	0.660	7
50 000	14.0	0.453	0.666	0.679	11
100 000	32.2	0.456	0.667	0.680	11

6 Conclusion

The experiments show that on commodity hardware SVMs can be used to train with millions of samples achieving at least decent errors in a few hours—even including 5-fold cross validation on a 10×10 hyper-parameter grid. The results for ECBDL demonstrate that local SVMs scale trivially across clusters. Hence, cloud scaling of fully automatic state-of the art SVMs is possible. Together with the statistical guarantees in Theorem 3.2 it is clear that local SVMs provide a reliable and broadly usable machine learning system for large scale data.

Acknowledgments

We thank the Institute of Mathematics at the University of Zurich for the computational resources.

References

- [1] K. P. Bennett and J. A. Blue. A support vector machine approach to decision trees. In *Neural Networks Proceedings, 1998. IEEE World Congress on Computational Intelligence. The 1998 IEEE International Joint Conference on*, volume 3, pages 2396–2401, 1998.
- [2] L. Bottou, O. Chapelle, D. DeCoste, and J. Weston. *Large-Scale Kernel Machines*. MIT Press, Cambridge, MA, 2007.
- [3] L. Bottou and V. Vapnik. Local Learning Algorithms. *Neural Computation*, 4:888–900, 1992.
- [4] B. Carl and I. Stephani. *Entropy, Compactness and the Approximation of Operators*. Cambridge University Press, Cambridge, 1990.
- [5] F. Chang, C.-Y. Guo, X.-R. Lin, and C.-J. Lu. Tree Decomposition for Large-Scale SVM Problems. *J. Mach. Learn. Res.*, 11:2935–2972, 2010.
- [6] H. Cheng, P.-N. Tan, and R. Jin. Localized Support Vector Machine and Its Efficient Algorithm. In *SIAM International Conference on Data Mining*, 2007.
- [7] H. Cheng, P.-N. Tan, and R. Jin. Efficient Algorithm for Localized Support Vector Machine. *IEEE Transactions on Knowledge and Data Engineering*, 22:537–549, 2010.
- [8] M. Eberts and I. Steinwart. Optimal learning rates for least squares SVMs using Gaussian kernels. In J. Shawe-Taylor, R.S. Zemel, P. Bartlett, F.C.N. Pereira, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24*, pages 1539–1547, 2011.
- [9] M. Eberts and I. Steinwart. Optimal learning rates for localized SVMs. Technical report, <http://arxiv.org/pdf/1507.06615v1>, 2015.
- [10] R. Hable. Universal consistency of localized versions of regularized kernel methods. *J. Mach. Learn. Res.*, 14:153–186, 2013.
- [11] J. Lin, L. Rosasco, and D.-X. Zhou. Iterative regularization for learning with convex loss functions. *arXiv:1503.08985*, 2015.

- [12] J. Platt. Fast training of support vector machines using sequential minimal optimization. In *Advances in Kernel Methods–Support Vector Learning*, pages 185–208. MIT Press, Cambridge, MA, 1999.
- [13] A. Rahimi and B. Recht. Random features for large-scale kernel machines. In J.C. Platt, D. Koller, Y. Singer, and S.T. Roweis, editors, *Advances in Neural Information Processing Systems 20*, pages 1177–1184. 2008.
- [14] L. Rosasco and S. Villa. Learning with incremental iterative regularization. <http://arxiv.org/pdf/1405.0042v2>, 2015.
- [15] B. Sriperumbudur and Z. Szabo. Optimal rates for random fourier features. In C. Cortes, N.D. Lawrence, D.D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 1144–1152. 2015.
- [16] I. Steinwart and A. Christmann. *Support Vector Machines*. Springer, New York, 2008.
- [17] I. Steinwart and C. Scovel. Fast rates for support vector machines using Gaussian kernels. *Ann. Statist.*, 35:575–607, 2007.
- [18] A. B. Tsybakov. Optimal aggregation of classifiers in statistical learning. *Ann. Statist.*, 32:135–166, 2004.
- [19] V. Vapnik and L. Bottou. Local Algorithms for Pattern Recognition and Dependencies Estimation. *Neural Computation*, 5:893–909, 1993.
- [20] C K. I. Williams and M. Seeger. Using the Nyström method to speed up kernel machines. In T.K. Leen, T.G. Dietterich, and V. Tresp, editors, *Advances in Neural Information Processing Systems 13*, pages 682–688. MIT Press, 2001.
- [21] D. Wu, K. P. Bennett, N. Cristianini, and J. Shawe-Taylor. Large Margin Trees for Induction and Transduction. In *Proceedings of the Sixteenth International Conference on Machine Learning*, ICML ’99, pages 474–483, San Francisco, CA, USA, 1999. Morgan Kaufmann Publishers Inc.
- [22] H. Zhang, A. C. Berg, M. Maire, and J. Malik. SVM-KNN: Discriminative Nearest Neighbor Classification for Visual Category Recognition. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 2, pages 2126–2136, 2006.
- [23] Y. Zhang, J. Duchi, and M. Wainwright. Divide and Conquer Kernel Ridge Regression: A Distributed Algorithm with Minimax Optimal Rates. *J. Mach. Learn. Res.*, 16:3299–3340, 2015.

A Proofs

A key concept to derive oracle inequalities and learning rates, which is used in the proof of Theorem 3.1 is the concept of entropy numbers, see [4] or [16, Definition A.5.26]. Recall that, for normed spaces $(E, \|\cdot\|_E)$ and $(F, \|\cdot\|_F)$ as well as an integer $i \geq 1$, the i -th (dyadic) entropy number of a bounded, linear operator $S : E \rightarrow F$ is defined by

$$e_i(S : E \rightarrow F) := e_i(SB_E, \|\cdot\|_F) \\ := \inf \left\{ \varepsilon > 0 : \exists s_1, \dots, s_{2^{i-1}} \in SB_E \text{ such that } SB_E \subset \bigcup_{j=1}^{2^{i-1}} (s_j + \varepsilon B_F) \right\},$$

where we use the convention $\inf \emptyset := \infty$, and B_E as well as B_F denote the closed unit balls in E and F , respectively.

Proof of Theorem 3.1. We denote by \tilde{H} the RKHS over X with Gaussian kernel of width γ_{\max} . Let $f_0 \in \tilde{H}$. Then, we can w.l.o.g. assume that $\|f_0\|_\infty \leq 1$, since the Gaussian kernel is bounded. For every $j \in \{1, \dots, m\}$ we define $f_j = \mathbf{1}_{A_j} f_0 = \widehat{f_0|_{A_j}}$ and remark that $f_j|_{A_j} \in H_{\gamma_{\max}}(A_j)$ due to [16, Exercise 4.4i)]. Hence, $f_j \in \hat{H}_{\gamma_{\max}}$ by definition of $\hat{H}_{\gamma_{\max}}$. Furthermore, since $\gamma_j \leq \gamma_{\max}$ for every $j \in \{1, \dots, m\}$, [16, Proposition 4.4.6] shows that $\hat{H}_{\gamma_{\max}} \subset \hat{H}_{\gamma_j}$ with

$$\|f_j\|_{\hat{H}_{\gamma_j}} \leq \left(\frac{\gamma_{\max}}{\gamma_j} \right)^{d/2} \|f_j\|_{\hat{H}_{\gamma_{\max}}}. \quad (5)$$

Hence, we find that $f_j \in \hat{H}_{\gamma_j}$ for every $j \in \{1, \dots, m\}$. Since

$$f_0 = \sum_{i=1}^m f_j$$

we conclude that $f_0 \in H$ by definition of H . Next, we observe with (5) that

$$\sum_{j=1}^m \lambda_j \|\mathbf{1}_{A_j} f_0\|_{\hat{H}_{\gamma_j}}^2 = \sum_{j=1}^m \lambda_j \|f_j\|_{\hat{H}_{\gamma_j}}^2 \leq \sum_{j=1}^m \lambda_j \left(\frac{\gamma_{\max}}{\gamma_j} \right)^d \|f_j\|_{\hat{H}}^2 \leq \sum_{j=1}^m \lambda_j \left(\frac{\gamma_{\max}}{\gamma_j} \right)^d \|f_0\|_{\hat{H}}^2.$$

By using the latter inequality and a bound for the approximation error, given in [16, Theorem 8.18], with tail exponent $\tau = \infty$ since X is compact and with $\lambda = \sum_{j=1}^m \lambda_j \left(\frac{\gamma_{\max}}{\gamma_j} \right)^d$, we find that

$$\begin{aligned} \sum_{j=1}^m \lambda_j \|\mathbf{1}_{A_j} f_0\|_{\hat{H}_{\gamma_j}}^2 + \mathcal{R}_{L,P}(f_0) - \mathcal{R}_{L,P}^* &\leq \sum_{j=1}^m \lambda_j \left(\frac{\gamma_{\max}}{\gamma_j} \right)^d \|f_0\|_{\hat{H}}^2 + \mathcal{R}_{L,P}(f_0) - \mathcal{R}_{L,P}^* \\ &\leq \max\{c_d, \tilde{c}_{d,\beta} c_{\text{NE}}\} \left(\sum_{j=1}^m \lambda_j \left(\frac{\gamma_{\max}}{\gamma_j} \right)^d \gamma_{\max}^{-d} + \gamma_{\max}^\beta \right) \\ &\leq \hat{c} \left(\sum_{j=1}^m \lambda_j \gamma_j^{-d} + \gamma_{\max}^\beta \right) \end{aligned} \quad (6)$$

where $\hat{c} := \max\{c_d, \tilde{c}_{d,\beta} c_{\text{NE}}\}$ with $c_d, \tilde{c}_{d,\beta} > 0$. Next, Theorem [9, Theorem 6] provides the bound $e_i(\text{id} : H_\gamma(A_j) \rightarrow L_2(\mathbb{P}_{X|A_j})) \leq a_j i^{-\frac{1}{2p}}$ for $i \geq 1$ with $a_j = \tilde{c}_p \sqrt{\mathbb{P}_X(A_j)} r^{\frac{d+2p}{2p}} \gamma_j^{-\frac{d+2p}{2p}}$, where \tilde{c}_p is a positive constant depending from p . For the constant a from Theorem B.1 this yields

$$\left(\max \left\{ c_p \sqrt{m} \left(\sum_{j=1}^m \lambda_j^{-p} a_j^{2p} \right)^{\frac{1}{2p}}, 2 \right\} \right)^{2p}$$

$$\begin{aligned}
&= \left(\max \left\{ c_p \sqrt{m} \left(\sum_{j=1}^m \lambda_j^{-p} \left(\tilde{c}_p \sqrt{P_X(A_j)} r^{\frac{d+2p}{2p}} \gamma_j^{-\frac{d+2p}{2p}} \right)^{2p} \right)^{\frac{1}{2p}}, 2 \right\} \right)^{2p} \\
&= \left(\max \left\{ c_p \tilde{c}_p m^{\frac{1}{2}} r^{\frac{d+2p}{2p}} \left(\sum_{j=1}^m \left(\lambda_j^{-1} \gamma_j^{-\frac{d+2p}{p}} P_X(A_j) \right)^p \right)^{\frac{1}{2p}}, 2 \right\} \right)^{2p} \\
&\leq \left(\max \left\{ c_p \tilde{c}_p m^{\frac{1}{2p}} r^{\frac{d+2p}{2p}} \left(\sum_{j=1}^m \lambda_j^{-1} \gamma_j^{-\frac{d+2p}{p}} P_X(A_j) \right)^{\frac{1}{2}}, 2 \right\} \right)^{2p} \\
&\leq \left(\max \left\{ c_p \tilde{c}_p 3^{\frac{d}{2p}} r \left(\sum_{j=1}^m \lambda_j^{-1} \gamma_j^{-\frac{d+2p}{p}} P_X(A_j) \right)^{\frac{1}{2}}, 2 \right\} \right)^{2p} \\
&\leq C_p r^{2p} \left(\sum_{j=1}^m \lambda_j^{-1} \gamma_j^{-\frac{d+2p}{p}} P_X(A_j) \right)^p + 4^p \\
&=: a^{2p},
\end{aligned}$$

where we used that $\|\cdot\|_p \leq m^{\frac{1-p}{p}} \|\cdot\|_1$ for $0 < p < 1$, as well as $mr^d \leq 16^d$ by (3) and that $C_p := c_p^{2p} \tilde{c}_p^{2p} 16^d$. Then, by using Theorem B.1, (6), the concavity of the function $t \mapsto t^{\frac{q+1}{q+2-p}}$ for $t \geq 0$ and the fact that $\tau \geq 1$ with $\tau \leq n$ we obtain that

$$\begin{aligned}
&\sum_{j=1}^m \lambda_j \|f_{D_j, \lambda_j, \gamma_j}\|_{\hat{H}_j}^2 + \mathcal{R}_{L,P}(\hat{f}_{D, \lambda, \gamma}) - \mathcal{R}_{L,P}^* \\
&\leq 9 \left(\sum_{j=1}^m \lambda_j \|\mathbf{1}_{A_j} f_0\|_{\hat{H}_{\gamma_j}}^2 + \mathcal{R}_{L,P}(f_0) - \mathcal{R}_{L,P}^* \right) + C \left(a^{2p} n^{-1} \right)^{\frac{q+1}{q+2-p}} + 3 \left(\frac{432 c_{\text{NE}}^{\frac{q}{q+1}} \tau}{n} \right)^{\frac{q+1}{q+2}} + \frac{30\tau}{n} \\
&\leq 9\hat{c} \left(\sum_{j=1}^m \lambda_j \gamma_j^{-d} + \gamma_{\max}^\beta \right) + C \left(C_p r^{2p} \left(\sum_{j=1}^m \lambda_j^{-1} \gamma_j^{-\frac{d+2p}{p}} P_X(A_j) \right)^p n^{-1} + 4^p n^{-1} \right)^{\frac{q+1}{q+2-p}} \\
&\quad + 3 \left(\frac{432 c_{\text{NE}}^{\frac{q}{q+1}} \tau}{n} \right)^{\frac{q+1}{q+2}} + \frac{30\tau}{n} \\
&\leq 9\hat{c} \left(\sum_{j=1}^m \lambda_j \gamma_j^{-d} + \gamma_{\max}^\beta \right) + C \left(C_p r^{2p} \left(\sum_{j=1}^m \lambda_j^{-1} \gamma_j^{-\frac{d+2p}{p}} P_X(A_j) \right)^p n^{-1} \right)^{\frac{q+1}{q+2-p}} \\
&\quad + C \left(\frac{4^p \tau}{n} \right)^{\frac{q+1}{q+2-p}} + 3 \left(\frac{432 c_{\text{NE}}^{\frac{q}{q+1}} \tau}{n} \right)^{\frac{q+1}{q+2}} + \frac{30\tau}{n} \\
&\leq \tilde{C}_{\beta, d, p, q} \left(\sum_{j=1}^m \lambda_j \gamma_j^{-d} + \gamma_{\max}^\beta + \left(r^{2p} \left(\sum_{j=1}^m \lambda_j^{-1} \gamma_j^{-\frac{d+2p}{p}} P_X(A_j) \right)^p n^{-1} \right)^{\frac{q+1}{q+2-p}} \right) \\
&\quad + \left(\frac{\tau}{n} \right)^{\frac{q+1}{q+2-p}} + \left(\frac{\tau}{n} \right)^{\frac{q+1}{q+2}} + \frac{\tau}{n} \\
&\leq C_{\beta, d, p, q} \left(\sum_{j=1}^m \lambda_j \gamma_j^{-d} + \gamma_{\max}^\beta + \left(r^{2p} \left(\sum_{j=1}^m \lambda_j^{-1} \gamma_j^{-\frac{d+2p}{p}} P_X(A_j) \right)^p n^{-1} \right)^{\frac{q+1}{q+2-p}} + \left(\frac{\tau}{n} \right)^{\frac{q+1}{q+2}} \right)
\end{aligned}$$

holds with probability P^n not less than $1 - 3e^{-\tau}$, where the constants $\tilde{C}_{\beta, d, p, q}$ and $C_{\beta, d, p, q}$ are

given by $\tilde{C}_{\beta,d,p,q} := \max\{9\hat{c}, C \cdot C_p^{\frac{q+1}{q+2-p}}, 3C_4^{\frac{p(q+1)}{q+2-p}}, 3c_{\text{NE}}^{\frac{q}{q+2}} \cdot 432^{\frac{q+1}{q+2}}, 30\}$ and $C_{\beta,d,p,q} := \max\{9\hat{c}, C \cdot C_p^{\frac{q+1}{q+2-p}}, 3C_4^{\frac{p(q+1)}{q+2-p}}, 9c_{\text{NE}}^{\frac{q}{q+2}} \cdot 432^{\frac{q+1}{q+2}}, 90\}$. \square

Proof of Theorem 3.2. First we simplify the presentation by using the sequences $\tilde{\lambda}_n := c_2 n^{-(\beta+d)\kappa}$ and $\tilde{\gamma}_n := c_3 n^{-\kappa}$ with $\kappa := \frac{(q+1)}{\beta(q+2)+d(q+1)}$. Then, we find with Theorem 3.1 together with $r_n = c_1 n^{-\nu}$, $\lambda_{n,j} = r_n^d \tilde{\lambda}_n$ and $\gamma_{n,j} = \tilde{\gamma}_n$ and with $\sum_{j=1}^{m_n} P_X(A_j) = 1$ and $m_n \leq 16^d r_n^{-d}$ that

$$\begin{aligned}
& \mathcal{R}_{L,P}(\hat{f}_{D,\lambda,\gamma}) - \mathcal{R}_{L,P}^* \\
& \leq C_{\beta,d,p,q} \left(\sum_{j=1}^m \lambda_{n,j} \gamma_{n,j}^{-d} + \gamma_{\max}^\beta + \left(r_n^{2p} \left(\sum_{j=1}^m \lambda_{n,j}^{-1} \gamma_{n,j}^{-\frac{d+2p}{p}} P_X(A_j) \right)^p n^{-1} \right)^{\frac{q+1}{q+2-p}} + \left(\frac{\tau}{n} \right)^{\frac{q+1}{q+2}} \right) \\
& = C_{\beta,d,p,q} \left(m r_n^d \tilde{\lambda}_n \tilde{\gamma}_n^{-d} + \tilde{\gamma}_n^\beta + \left(r_n^{2p} \left(r_n^{-d} \tilde{\lambda}_n^{-1} \tilde{\gamma}_n^{-\frac{d+2p}{p}} \sum_{j=1}^m P_X(A_j) \right)^p n^{-1} \right)^{\frac{q+1}{q+2-p}} + \left(\frac{\tau}{n} \right)^{\frac{q+1}{q+2}} \right) \\
& \leq 3^d C_{\beta,d,p,q} \left(\tilde{\lambda}_n \tilde{\gamma}_n^{-d} + \tilde{\gamma}_n^\beta + \left(r_n^{p(2-d)} \tilde{\lambda}_n^{-p} \tilde{\gamma}_n^{-(d+2p)} n^{-1} \right)^{\frac{q+1}{q+2-p}} + \left(\frac{\tau}{n} \right)^{\frac{q+1}{q+2}} \right) \\
& \leq C_{\beta,\nu,d,p,q} \left(n^{-(\beta+d)\kappa} n^{+d\kappa} + n^{-\beta\kappa} + \left(\frac{n^{-p\nu(2-d)}}{n^{-p(\beta+d)\kappa} n^{-(d+2p)\kappa+1}} \right)^{\frac{q+1}{q+2-p}} + \left(\frac{\tau}{n} \right)^{\frac{q+1}{q+2}} \right) \\
& = C_{\beta,\nu,d,p,q} \left(2n^{-\beta\kappa} + \left(\frac{n^{-p[\nu(2-d) - (\beta+d)\kappa - 2\kappa]}}{n^{-d\kappa+1}} \right)^{\frac{q+1}{q+2-p}} + \left(\frac{\tau}{n} \right)^{\frac{q+1}{q+2}} \right) \\
& = C_{\beta,\nu,d,p,q} \left(2n^{-\beta\kappa} + \frac{n^{-\frac{p(q+1)}{q+2-p}[\nu(2-d) - (\beta+d+2)\kappa]}}{n^{\frac{\beta(q+2)(q+1)}{(\beta(q+2)+d(q+1))(q+2-p)}}} + \left(\frac{\tau}{n} \right)^{\frac{q+1}{q+2}} \right) \\
& \leq C_{\beta,\nu,d,p,q} \left(2n^{-\beta\kappa} + \frac{n^{-\frac{p(q+1)}{q+2-p}[\nu(2-d) - (\beta+d+2)\kappa]}}{n^{\frac{\beta(q+1)}{(\beta(q+2)+d(q+1))}}} + \left(\frac{\tau}{n} \right)^{\frac{q+1}{q+2}} \right) \\
& \leq C_{\beta,\nu,d,p,q} \left(2n^{-\frac{\beta(q+1)}{\beta(q+2)+d(q+1)}} + n^{-\frac{\beta(q+1)}{\beta(q+2)+d(q+1)}} n^{-\frac{p(q+1)}{q+2}[\nu(2-d) - \frac{(\beta+d+2)(q+1)}{\beta(q+2)+d(q+1)}]} + (\tau n^{-1})^{\frac{q+1}{q+2}} \right) \\
& \leq C_{\beta,\nu,\xi,d,q} \tau^{\frac{q+1}{q+2}} \cdot n^{-\frac{\beta(q+1)}{\beta(q+2)+d(q+1)}} + \xi
\end{aligned}$$

holds with probability P^n not less than $1 - 3e^{-\tau}$, where the constants $C_{\beta,\nu,d,p,q}, c_1, c_2, c_3 > 0$ depend on β, ν, d, p, q and the constant $C_{\beta,\nu,\xi,d,q} > 0$ depends on β, ν, ξ, d, q . Furthermore, we remark that we chose p sufficiently close to zero such that $\xi \geq \frac{p(q+1)}{q+2} \left(\nu(d-2) + \frac{(\beta+d+2)(q+1)}{\beta(q+2)+d(q+1)} \right) \geq 0$. \square

B Appendix

Theorem B.1. Let P be a distribution on $X \times Y$ with noise exponent $q \in (0, \infty]$ and let $L : Y \times \mathbb{R} \rightarrow [0, \infty]$ be the hinge loss. Furthermore let (A) be satisfied with $H_j := H_{\gamma_j}(A_j)$ and assume that, for fixed $n \geq 1$, there exist constants $p \in (0, 1)$ and $a_1, \dots, a_m > 0$ such that for all $j \in \{1, \dots, m\}$

$$e_i(\text{id} : H_j \rightarrow L_2(P_{X|A_j})) \leq a_j i^{-\frac{1}{2p}}, \quad i \geq 1. \quad (7)$$

Finally, fix an $f_0 \in H$. Then, for all fixed $\tau > 0$, $\lambda = (\lambda_1, \dots, \lambda_m) > 0$, $\gamma = (\gamma_1, \dots, \gamma_m) > 0$ and

$$a := \max \left\{ c_p \sqrt{m} \left(\sum_{j=1}^m \lambda_j^{-p} a_j^{2p} \right)^{\frac{1}{2p}}, 2 \right\}$$

the VP-SVM given by (4) satisfies

$$\begin{aligned} & \sum_{j=1}^m \lambda_j \|f_{D_j, \lambda_j, \gamma_j}\|_{\hat{H}_j}^2 + \mathcal{R}_{L, P}(\hat{f}_{D, \lambda, \gamma}) - \mathcal{R}_{L, P}^* \\ & \leq 9 \left(\sum_{j=1}^m \lambda_j \|\mathbf{1}_{A_j} f_0\|_{\hat{H}_j}^2 + \mathcal{R}_{L, P}(f_0) - \mathcal{R}_{L, P}^* \right) + C (a^{2p} n^{-1})^{\frac{q+1}{q+2-p}} + 3 \left(\frac{432 c_{NE}^{\frac{q}{q+1}} \tau}{n} \right)^{\frac{q+1}{q+2}} + \frac{30\tau}{n} \end{aligned}$$

with probability P^n not less than $1 - 3e^{-\tau}$, where $C > 0$ is a constant only depending on p .

Proof of Theorem B.1. The result follows from a direct application of [9, Theorem 5]. To this end, we note that the hinge loss is Lipschitz continuous and can be clipped at $M = 1$. Since H is the sum-RKHS of RKHSs with Gaussian kernels and the Gaussian kernel is bounded, w.l.o.g. we assume for $f_0 \in H$ that $\|f_0\|_\infty \leq 1$. Hence, $\|L \circ f_0\|_\infty \leq 2$ and therefore $B_0 = 2$. Furthermore [16, Theorem 8.24] showed that for the hinge loss the constants V and ϑ from [9, Theorem 5] can be achieved by $V = 6c_{NE}^{\frac{q}{q+1}}$ and $\vartheta = \frac{q}{q+1}$. That means

$$\begin{aligned} & \sum_{j=1}^m \lambda_j \|f_{D_j, \lambda_j, \gamma_j}\|_{\hat{H}_j}^2 + \mathcal{R}_{L, P}(\hat{f}_{D, \lambda, \gamma}) - \mathcal{R}_{L, P}^* \\ & \leq 9 \left(\sum_{j=1}^m \lambda_j \|\mathbf{1}_{A_j} f_0\|_{\hat{H}_j}^2 + \mathcal{R}_{L, P}(f_0) - \mathcal{R}_{L, P}^* \right) + C \left(\frac{a^{2p}}{n} \right)^{\frac{1}{2-p-\vartheta+p}} + 3 \left(\frac{72V\tau}{n} \right)^{\frac{1}{2-\vartheta}} + \frac{15B_0\tau}{n} \\ & = 9 \left(\sum_{j=1}^m \lambda_j \|\mathbf{1}_{A_j} f_0\|_{\hat{H}_j}^2 + \mathcal{R}_{L, P}(f_0) - \mathcal{R}_{L, P}^* \right) + C \left(\frac{a^{2p}}{n} \right)^{\frac{q+1}{q+2-p}} + 3 \left(\frac{432 c_{NE}^{\frac{q}{q+1}} \tau}{n} \right)^{\frac{q+1}{q+2}} + \frac{30\tau}{n} \end{aligned}$$

holds with probability P^n not less than $1 - 3e^{-\tau}$. \square

C Some more details for results

Table 6: Training times divided by the product of training set size, cell size and dimension (in seconds times 10^9). This shows that the time complexity in (1) is fulfilled quite nicely.

	2000	5000	10000	15000
	training time			
HIGGS	33	29	29	29
HEPMASS	22	19	19	20
GASSENSOR	20	19	19	19
SUSY	45	39	38	38
COVTYPE	10	9	9	9
COD-RNA	52	51	45	48
SKIN	116	117	106	111




Figure 4: Training time vs. test error, this is the final trade-off.